

Estimation of covariance matrices based on hierarchical inverse-Wishart priors[☆]

M. Bouriga^{a,b,*}, O. Féron^{a,b}

^a*Université Paris IX - Dauphine, Place du Maréchal de Lattre de Tassigny, 75116 Paris, France*

^b*EDF R&D, 1 avenue Général de Gaulle, 92140 Clamart, France*

Abstract

This paper focuses on Bayesian shrinkage for covariance matrix estimation. We examine posterior properties and frequentist risks of Bayesian estimators based on new hierarchical inverse-Wishart priors. More precisely, we give the existence conditions of the posterior distributions. Advantages in terms of numerical simulations of posteriors are shown. A simulation study illustrates the performance of the estimation procedures under three loss functions for relevant sample sizes and various covariance structures.

Keywords: Bayesian covariance estimation; Shrinkage; Hierarchical Inverse-Wishart prior; Loss function; Experiments comparisons.

1. Introduction

Estimating a covariance matrix efficiently is an important statistical issue. Often, applied scientific problems require an estimate of a covariance matrix in the context of a large matrix dimension p relative to the number of observations n . In such settings, standard estimators - the sample covariance matrix or the maximum likelihood estimator - are known to perform poorly [26, 27, 4]. When n is smaller than p , they are not positive definite. When it is larger, they are invertible but still inappropriate because unstable, unless $\frac{p}{n}$ is negligibly small. Indeed, if n is of the same order as p , the sample eigenvalues significantly deviate from the population eigenvalues [27, 4]. This fact has incited many authors to focus on the eigenvalues attempting to overcome their distortion.

Some approaches to more stably estimating the matrix in small samples have been proposed, such as shrinkage methods on which we will focus. The work along these directions can be found in both frequentist and Bayesian frameworks. The proposed estimators for the covariance matrix are then derived from

[☆]This document is a part of the thesis of Mathilde Bouriga.

*Corresponding author. Fax: 0147653900.

Email addresses: mathilde.bouriga@edf.fr (M. Bouriga), olivier-2.feron@edf.fr (O. Féron)

a decision-theoretic perspective or associated with an appropriate prior and a specific loss function.

James and Stein [14] were the first to propose biased estimators for covariance matrix under Stein's loss function, dominating the classical sample covariance matrix. Since, many authors have explored improved James-Stein type estimators under Stein's loss [27, 5, 17, 20] or other losses [8, 12, 16, 20]. Ledoit and Wolf [18] consider Steinian shrinkage toward the single-index covariance matrix to estimate the covariance matrix of stock returns. Furrer and Bengtsson [10] consider "tapering" the sample covariance matrix, that is, gradually shrinking the off-diagonal elements toward zero.

From a Bayesian perspective, the common approach [13, 2] yielding estimators which shrink towards a structure uses the conjugate prior on the covariance matrix, an inverse-Wishart distribution with some degrees of freedom and a scale matrix as hyperparameters. The appeal of conjugate priors is to allow efficient posterior simulations but such priors might be contested because of their lack of flexibility; specifically, with an inverse-Wishart prior, only one parameter does control the variability of the matrix elements. Efficient Bayesian estimators involving more flexible priors are obtained using various decompositions, the most well-known are derived from the variance-correlation strategy introduced by Barnard and al. [1, 3], the spectral [3, 28] or the Cholesky [25] decompositions of the covariance matrix and the matrix-logarithmic covariance model [19]. In particular, in [1], the covariance matrix is modeled in terms of standard deviations and correlations. After discussion about suitable priors, the authors advocate a flat prior on the space of correlation matrices and log normal priors on the variances. Yang and Berger [28] develop a reference prior approach for the covariance matrix, approach known to outperform another common noninformative prior, Jeffreys prior [15], for high-dimensional problems. Working with the spectral decomposition of the matrix, their method shrinks the eigenvalues, it results in a better estimation of the underlying eigenstructure. Smith and Kohn [25] use a prior that allows for zero elements in the strict lower triangle of the Cholesky decomposition of the inverse of the covariance matrix. Leonard and Hsu [19] model the matrix logarithm of the covariance matrix and place a multivariate normal distribution on the $\frac{p(p+1)}{2}$ vectorized non-redundant elements of this matrix. Daniels and Kass [3] focus on shrinking the matrix toward a structure, specifically, a diagonal matrix, using a fully Bayesian approach and using three different priors: normal priors on the z-transform of the correlations, normal priors on the logit of the Givens angles and inverse-Wishart prior. These parametrizations do not have simple statistical interpretation and the use of these methods has been limited in view of the difficulties of computation implied.

Moreover all these parametrizations do not solve the difficulty in handling and estimating hyperparameters. Empirical Bayes estimates of the hyperparameters can be proposed, for instance by [13, 2, 24] in the case of inverse-Wishart prior. An alternative to this empirical specification is to use hierarchical modeling. Multilevel modeling will ensure that uncertainty in higher level param-

eters propagates into inferences on lower level parameters and is supposed to be more flexible and also more stable than those based on diffuse priors [11]. Applied to the case of inverse-Wishart prior, it is intended to offer objectivity in terms of how close the true matrix is to the specified structure by allowing for data-dependent shrinkage towards this structure. To our knowledge, hierarchical modelization based on inverse-Wishart priors has only been proposed in [3] with a constraint of fixing an upper bound on the degrees of freedom to ensure proper posterior distributions.

In this paper we focus on hierarchical inverse-Wishart priors on the covariance matrix such that shrinkage toward diagonality is involved. This approach builds on the works of Daniels and Kass [3]. After relaxing the prior on the degrees of freedom, we establish the precise conditions to ensure the properness of the posterior distributions. We also give a detailed experimental comparison of the competing Bayesian models and the maximum likelihood estimator under different loss functions. According to the loss, the priors don't have the same effect on posterior inference and we outline the limits of the hierarchical inverse-Wishart priors as a "default" choice for the covariance matrix estimation.

The paper is structured as follows. Section 2 presents the Normal inverse-Wishart model as covariance shrinkage modeling approach. In Section 3 we define the three hierarchical priors that we will consider and conditions to get proper posterior distributions. Section 4 describes the attractive Markov Chain Monte Carlo sampling scheme for Bayesian computation of posteriors. Section 5 reports numerical results for the matrix estimators and for the eigenvalues estimators. The last section presents some conclusions and provides recommendations for using such priors.

2. The normal inverse-Wishart model

Let $\mathbf{X} = (X^{(1)}, \dots, X^{(p)})^T$ be a p -dimensional random vector following a multivariate normal distribution $\mathcal{N}_p(0, \Sigma)$. Σ is an unknown covariance matrix and belongs to the set of $p \times p$ symmetric positive definite matrices \mathcal{S}^+ . Given an independent and identically distributed sample $(\mathbf{X}_1, \dots, \mathbf{X}_n)$ of \mathbf{X} , we wish to estimate the covariance matrix. The associated likelihood function for Σ is

$$L(\Sigma|S) = \frac{|\Sigma|^{-\frac{n}{2}}}{(2\pi)^{\frac{np}{2}}} \exp \left\{ -\frac{1}{2} \text{tr} (\Sigma^{-1}S) \right\} \quad (1)$$

where $S = \sum_{i=1}^n \mathbf{X}_i^T \mathbf{X}_i$ is the scatter matrix.

The maximum likelihood estimator of Σ , $\hat{\Sigma}_{MLE}$, defined by $\frac{S}{n}$, is a classical estimator of Σ . However, it becomes unstable when p is moderate or large relative to the sample size n because of the large number - $\frac{p(p+1)}{2}$ - of unknown parameters to be estimated. When $n < p$, $\hat{\Sigma}_{MLE}$ is no longer positive definite. Even when $n > p$, the matrix S is positive definite but does lead to the distortion of the eigenstructure [27, 4] for p close to n , especially when the true matrix is close to be diagonal. That motivates the choice of a Bayesian regularization approach and leads to assign a prior on Σ .

In the absence of reliable prior information, the selection of the prior distribution is quite delicate and generic solutions must be chosen instead. Here we consider a general class of priors. Since the model \mathbf{X} when Σ^{-1} varies is from a natural exponential family of distributions, a commonly used class of distributions for the canonical parameter Σ^{-1} is the conjugate family as defined by Diaconis & Ylvisaker [6], called the Wishart distribution. The induced prior on Σ is then the inverse-Wishart.

In the notation of Eaton [7], $\Sigma|\beta, D \sim \mathcal{IW}(\beta, D)$ means that Σ has the inverse-Wishart distribution with degrees of freedom $\beta > p-1$ and scale matrix $D \in \mathcal{S}^+$. The density is given by

$$\pi(\Sigma|\beta, D) = \frac{D^{\frac{\beta}{2}}}{C(\beta, p)} |\Sigma|^{-\frac{\beta+p+1}{2}} \exp \left\{ -\frac{1}{2} \text{tr} (\Sigma^{-1} D) \right\}. \quad (2)$$

The normalising constant turns out to be equal to

$$C(\beta, p) = 2^{\frac{\beta p}{2}} \Gamma_p \left(\frac{\beta}{2} \right) \quad (3)$$

where $\Gamma_p(\cdot)$ is the multivariate Gamma function defined as $\Gamma_p(a) = \pi^{\frac{p(p-1)}{4}} \prod_{j=1}^p \Gamma(a + \frac{1-j}{2})$.

The restriction that β be greater than $p-1$ is necessary so that $\Gamma(\frac{\beta+1-j}{2})$ be well defined. Moreover, as $\mathbb{E}(\Sigma) = \frac{D}{\beta-p-1}$, the expectation of Σ will exist if and only if $\beta > p+1$.

Now, let $(\mathbf{X}_1, \dots, \mathbf{X}_n)$ be a Gaussian sample associated with an inverse-Wishart prior on Σ centered in D , the posterior mean of Σ , a likely estimator of Σ , is then equal to

$$\mathbb{E}(\Sigma|S) = \frac{(\beta-p-1)D+S}{\beta+n-p-1} = \frac{(\beta-p-1)D+n\hat{\Sigma}_{MLE}}{\beta+n-p-1} \quad \text{with } \beta > p+1.$$

It shows that β controls the amount of shrinkage: when n is held fixed and β is allowed to grow, the posterior mean tends towards D while the estimator tends towards $\hat{\Sigma}_{MLE}$ if β is held fixed and n is allowed to grow.

When we consider the eigenvalues of the posterior mean, it is easy to see that the eigenvalues g_i , $i = 1, \dots, p$ of $\mathbb{E}(\Sigma|S)$ are

$$g_i = \frac{(\beta-p-1)d_{ii} + nl_i}{\beta+n-p-1}, \quad \forall i = 1, \dots, p. \quad (4)$$

where l_i are the eigenvalues of $\hat{\Sigma}_{MLE}$.

We can check that, for $l_i < d_{ii}$, we always have $l_i < g_i$ and, for $l_i > 1$, we have $l_i > g_i$.

Notice that we can consider $\mathbb{E}(\Sigma^{-1}|S)^{-1}$ as the estimator of Σ , you will find the same kind of results.

In summary, the span of the eigenvalues of Bayes estimators based on inverse-Wishart prior will be smaller than the span of the eigenvalues of S , which could be used to correct the instability of the standard estimators.

3. Inverse-Wishart prior distribution: choice of the hyperparameters

As seen in Section 2 the inverse-Wishart distribution is characterized by the hyperparameters β and D . Sometimes the hyperparameters can be specified by the investigators but there is rarely good scientific information on which to base these specifications. They can be obtained by empirical Bayes estimation [13, 2, 24]. A frequently-applied procedure is to set the scale matrix D equal to $\hat{\Sigma}_{MLE}$ and p degrees of freedom. However it turns out to be not convenient as soon as $\hat{\Sigma}_{MLE}$ becomes suspect.

Consequently, we prefer to embed a structure for D . The most commonly employed matrix targets are the identity matrix and its scalar multiples [24]. They are low-dimensional, thus they impose a rather strong structure which in turn requires only little data to fit the hyperparameters remaining to estimate. Alternatively we can assign a further prior distribution on these hyperparameters. This implies hierarchical models which allow a more objective approach to inference [11]. They are supposed to provide more flexibility than non-hierarchical priors. As for the degrees of freedom β , they can either be taken as small as possible with the idea that large values supporting the scale matrix structure or they can be given their own prior distribution.

Here we adopt a fully Bayesian approach: we investigate three inverse-Wishart hierarchical priors with unknown degrees of freedom and scale matrix, inducing shrinkage towards diagonality.

3.1. Daniels and Kass prior

Daniels and Kass [3] assign flat improper priors on the logarithm of the elements of the diagonal scale matrix and a vague proper uniform distribution on the logarithm of the degrees of freedom, over $]p-1; b]$, with b a large value. Let A be a $p \times p$ diagonal definite-positive matrix, the Daniels and Kass model, namely Model D&K, is defined by

$$\begin{aligned}\Sigma|\alpha_j, \beta &\sim \mathcal{IW}(\beta, A), \quad A = \text{diag}(\alpha_1, \dots, \alpha_p) \\ \pi(\alpha_j) &\propto \frac{1}{\alpha_j} \mathbb{I}_{]0; +\infty[}(\alpha_j) \\ \pi(\beta) &\propto \frac{1}{\beta} \mathbb{I}_{]p-1; b]}(\beta).\end{aligned}\tag{5}$$

From (1) and (5), the joint posterior distribution given the data \mathbf{X} is equal to

$$\pi(\Sigma, A, \beta|S) = \frac{|\Sigma^{-1}|^{\frac{\beta+n+p+1}{2}} \exp\left[-\frac{1}{2}\text{tr}(\Sigma^{-1}(S+A))\right] |A|^{\frac{\beta}{2}-1}}{(2\pi)^{\frac{pn}{2}} 2^{\frac{p\beta}{2}} \beta \Gamma_p\left(\frac{\beta}{2}\right)} \mathbb{I}_{S^+ \times (\mathbb{R}_+^*)^p \times]p-1; b]}(\Sigma, \alpha_1, \dots, \alpha_p, \beta).\tag{6}$$

The bound b must be a finite value to keep the joint posterior distribution proper, with constraint to be large in order to minimize its effect on inference. Anyway we exclude arbitrarily some values of β : we never know how large is large enough. Another similar but improper prior on the logarithm of β is preferred, rather than a proper one with bounded support.

3.2. A "diagonal, equal variance" model as prior mean matrix

We choose to center the prior matrix in αI_p . By this way, the resulting estimators will shrink all elements of S . Model 1 will be defined by

$$\begin{aligned}\Sigma|\alpha, \beta &\sim \mathcal{IW}(\beta, (\beta - p - 1)\alpha I_p) \\ \pi(\alpha) &\propto \frac{1}{\alpha} \mathbb{I}_{]0; +\infty[}(\alpha) \\ \pi(\beta) &\propto \frac{1}{\beta^\delta} \mathbb{I}_{]p+1; +\infty[}(\beta) \quad , \text{ with } \delta \text{ positive integer.}\end{aligned}\tag{7}$$

After a reparametrization given by $\phi = \alpha(\beta - p - 1)$, we derive (8), the joint posterior distribution of (Σ, ϕ, β) given S :

$$\pi(\Sigma, \phi, \beta|S) = \frac{|\Sigma^{-1}|^{\frac{\beta+n+p+1}{2}} \exp\left[-\frac{1}{2}\text{tr}(\Sigma^{-1}(S + \phi I))\right] |\phi I|^{\frac{\beta}{2}}}{(2\pi)^{\frac{pn}{2}} 2^{\frac{p\beta}{2}} \beta^\delta \Gamma_p(\frac{\beta}{2}) \phi} \mathbb{I}_{S^+ \times \mathbb{R}_+^* \times]p+1; +\infty[}(\Sigma, \phi, \beta).\tag{8}$$

The introduction of δ permits the limit superior of β in Model D&K to be relaxed. The problem of parameter choice persists but it is hoped that this model influences less the posterior distribution.

The conditions ensuring properness of the posterior distribution are driven by δ . For $\delta = 1$, we know by [3] that (8) is improper.

Proposition 1: *The posterior distribution (8) is a proper probability density function for all $\delta > 1$:*

$$\int_{p+1}^{+\infty} \int_{]0; +\infty[} \int_{S^+} \pi(\Sigma, \phi, \beta|S) d\Sigma d\phi d\beta < \infty\tag{9}$$

This proof is deferred to Appendix A.

It leads to a remark about the posterior distribution of β .

Corollary: *The posterior marginal distribution of β , $\pi(\beta|S)$, has its m -th moment and higher ones defined if $m < \delta - 1$.*

Consequently, if β is no longer considered as a nuisance parameter but as a parameter of interest, samples from $\pi(\beta|\mathbf{X})$ must be used with caution. For instance, for $\delta < 3$, the mean of sample paths from $\pi(\beta|\mathbf{X})$ has no sense since the posterior mean of β will not exist.

3.3. A "diagonal, unequal variance" model as prior mean matrix

We can prefer to keep a less strong structure for the matrix target, as in Section 3.1, where we only shrink the off-diagonal elements of S . Moreover we assume that the prior mean matrix does exists, equal to an arbitrary diagonal

matrix A . Let A be a $p \times p$ diagonal definite-positive matrix, Model 2 will be defined by

$$\begin{aligned}\Sigma|\alpha_j, \beta &\sim \mathcal{IW}(\beta, (\beta - p - 1)A), \quad A = \text{diag}(\alpha_1, \dots, \alpha_p) \\ \pi(\alpha_j) &\propto \frac{1}{\alpha_j} \mathbb{I}_{]0; +\infty[}(\alpha_j) \\ \pi(\beta) &\propto \frac{1}{\beta^\delta} \mathbb{I}_{]p+1; +\infty[}(\beta).\end{aligned}\tag{10}$$

Proceed to the change of variable $\Phi = (\beta - p - 1)A$ and refer the diagonal matrix Φ by its diagonal elements, ϕ_i , i from 1 to p . The joint posterior distribution is then given by

$$\pi(\Sigma, \Phi, \beta|S) = \frac{|\Sigma^{-1}|^{\frac{\beta+n+p+1}{2}} \exp\left[-\frac{1}{2}\text{tr}(\Sigma^{-1}(S + \Phi))\right] |\Phi|^{\frac{\beta}{2}-1}}{(2\pi)^{\frac{pn}{2}} 2^{\frac{p\beta}{2}} \beta^\delta \Gamma_p(\frac{\beta}{2})} \mathbb{I}_{S^+ \times (\mathbb{R}_+^*)^p \times]p+1; +\infty[}(\Sigma, \phi_1, \dots, \phi_p, \beta).\tag{11}$$

Proposition 2: *The posterior distribution (11) is a proper probability density function for all $\delta > 1$:*

$$\int_{p+1}^{+\infty} \int_{]0; +\infty[^p} \int_{S^+} \pi(\Sigma, \Phi, \beta|S) d\Sigma d\phi_1 \dots d\phi_p d\beta < \infty\tag{12}$$

The proof of Proposition 2 is deferred to the Appendix Appendix A. Note that the posterior distribution (11) for Model 2 appears quasi-identical to the density (6) for Model D&K, although the prior hypotheses differ. The existence of the prior mean matrix is only taken into account in the lower bound on β in the posterior distribution. The second difference is in the power of β in the denominator.

4. Computation of Bayes estimators

Now we can derive the Bayes estimators related to the prior models of Sections 3.2 and 3.3, under two loss functions. We will use Markov Chain Monte Carlo (MCMC) simulations to estimate these posterior quantities numerically.

4.1. Loss functions and associated estimators

A question that naturally arises in various contexts in multivariate analysis and related topics is whether to estimate Σ or its inverse. We choose to focus on the estimation of Σ in this paper. This parameter has a natural and well understood interpretation in multivariate analysis and its direct estimation has numerous applications.

In the following, the Bayes estimators of Σ are calculated with respect to two common loss functions:

- the squared Frobenius loss function

$$L_2(\hat{\Sigma}, \Sigma) = \text{tr}(\hat{\Sigma} - \Sigma)^2$$

- Stein's loss function

$$L_1(\hat{\Sigma}, \Sigma) = \text{tr}(\hat{\Sigma}\Sigma^{-1}) - \log \det(\hat{\Sigma}\Sigma^{-1}) - p$$

The corresponding Bayes estimators for Σ are, respectively, $\hat{\Sigma}_2 = \mathbb{E}(\Sigma|S)$ and $\hat{\Sigma}_1 = \mathbb{E}(\Sigma^{-1}|S)^{-1}$.

The L_2 loss corresponds to the equivalent of the squared error loss function in a matrix setting. Thus $L_2(\hat{\Sigma}, \Sigma) = \sum_{i=1}^p \sum_{j=1}^p (\hat{\sigma}_{ij} - \sigma_{ij})^2$ is a natural quadratic measure of distance between the true (Σ) and inferred covariance matrix ($\hat{\Sigma}$). The L_1 loss was introduced by Stein [27] to estimate the multinormal covariance matrix and also called entropy loss. It results from evaluating the divergence of Kullback-Leibler, namely $\int p(x) \log \left\{ \frac{p(x)}{q(x)} \right\} dx$ for two Gaussian distributions with densities $p(x)$ and $q(x)$ defined by covariance matrices Σ and $\hat{\Sigma}$. This scale invariant loss function will penalize the relative estimation error of the small eigenvalues, as illustrated in Section 5. Various alternative losses have also been proposed in the literature [28].

Both estimators are approximated by using the MCMC sampling algorithm described in Section 4.2.

4.2. MCMC algorithm for sampling posterior distributions

Once we introduce hierarchical priors, the conjugate structure that typically makes Gibbs sampling so attractive is no longer ensured. Fortunately, since the inverse-Wishart distributions is also a conditionally-conjugate family, the full conditional distribution of Σ is still inverse-Wishart. Moreover the full conditional for the additional parameters Φ and β can be also simulated easily. Sampling from the target posterior will call on iterative resampling from inverse-Wishart, inverse-Gamma distributions and from the posterior distribution of β . The simulation of the latter is based on a Metropolis sampling scheme, described in Algorithm 3.

A finite-sample distribution from the posterior distribution (8) of (Σ, ϕ, β) from Model 1 can then be obtained by a systematic scan Metropolis-Hasting-within-Gibbs algorithm [23], described by Algorithm 1.

Then to obtain samples from the posterior distribution (11) of (Σ, Φ, β) from Model 2, you need to follow all steps of Algorithm 1 except Step 4 and instead, use Step 4_{bis} described in Algorithm 2.

4.3. The Metropolis sampling method for β

The density of $\beta^{(k)}|\Sigma^{(k)}, \phi^{(k)}, S$ is such that

$$\pi(\beta^{(k)}|\Sigma^{(k)}, \phi^{(k)}, S) \propto \exp \left[\beta^{(k)} C^{(k)} - \delta \log \beta^{(k)} - \log(\Gamma_p \left(\frac{\beta^{(k)}}{2} \right)) \right] \mathbb{I}_{]p+1; +\infty[}(\beta^{(k)})$$

with $C^{(k)} = \frac{\log |\Sigma^{(k)-1}| + \log |\Phi^{(k)}| - p \log 2}{2}$.

(13)

Algorithm 1 Metropolis-Hasting-within-Gibbs sampling scheme for the joint posterior (8)

- 1: Initialization with $k = 0$ and arbitrary values for ϕ_0 and Σ_0
 - 2: Increment $k = k + 1$
 - 3: Draw a sample $\Sigma^{(k)} | \phi^{(k-1)}, \beta^{(k-1)}, S \sim \mathcal{IW}(\beta^{(k-1)} + n, S + \phi^{(k-1)}I)$
 - 4: Draw a sample $\phi^{(k)} | \beta^{(k-1)}, \Sigma^{(k)}, S \sim \mathcal{G}\left(\frac{p\beta^{(k-1)}}{2}, \frac{tr(\Sigma^{(k)} - 1)}{2}\right)$
 - 5: Draw a sample $\beta^{(k)} \sim \pi(\beta^{(k)} | \Sigma^{(k)}, \phi^{(k)}, S)$, see Algorithm 3
 - 6: Return to 2 except in the case of stop criterion
-

Algorithm 2 Modification in Algorithm 1 for sampling from the joint posterior (11)

4_{bis} : For j from 1 to p , draw a sample $\phi_j^{(k)} | \beta^{(k-1)}, \Sigma^{(k)}, S \sim \mathcal{G}\left(\frac{p\beta^{(k-1)}}{2}, \frac{\Sigma_{jj}^{(k)} - 1}{2}\right)$

We describe a random-walk Metropolis algorithm to sample from the distribution of γ^k when $\gamma^{(k)} = \log(\beta^{(k)} - p - 1)$. The proposed algorithm is a Markov chain with a Gaussian random-walk centered in the value $\gamma^{(k-1)}$ as symmetric proposal density q at iteration k .

$$q(\gamma^{prop}) = \mathcal{N}_p\left(\gamma^{(k-1)}, 2 \times \hat{\sigma}^{2(k-1)}\right).$$

$\hat{\sigma}^{2(k-1)}$ corresponds to the variance of γ^k , estimated by numerical integration. The acceptance probability is equal to $\min\left(1, \frac{q(\gamma^{prop})}{q(\gamma^{(k-1)})}\right)$. Samples from (13) can be obtained using the mapping described in Algorithm 3.

5. Simulation results

The objective of this section is to give a detailed comparison between different estimators: the proposed Bayes estimators and the maximum likelihood estimator. More specifically, the study will consider the true matrices used in [3], with different structures and conditionings¹. We will derive frequentist characteristics of the Bayesian procedures in terms of risks associated to three loss

¹A matrix is said ill-conditioned if the ratio of its maximum and minimum eigenvalue is large. The closer it is to 1, the better conditioned the matrix is.

Algorithm 3 The random-walk Metropolis algorithm for the joint posterior (13)

- 1: Set $\gamma^{(k-1)} = \log(\beta^{(k-1)} - p - 1)$
 - 2: Compute $\hat{\sigma}^{2(k-1)}$ by numerical integration
 - 3: Sample γ^{prop} from $q(\gamma^{prop} | \gamma^{(k-1)}, 2 \times \hat{\sigma}^{2(k-1)})$
 - 4: Set $\log(\rho) = \log q(\gamma^{prop}) - \log q(\gamma^{(k-1)})$
 - 5: Sample u from $\mathcal{U}_{[0;1]}$
 - 6: If $\log \rho > \log u$ then $\gamma^{(k)} = \gamma^{prop}$. Otherwise $\gamma^{(k)} = \gamma^{(k-1)}$
 - 7: Set $\beta^{(k)} = \exp(\gamma^{(k)}) + p + 1$
-

functions. Through this simulation study, we will get ideas about the behaviour of the competing estimation methods in different situations.

5.1. Simulation design

We carried out the simulation study from Section 3 of Daniels and Kass's paper [3]. They consider seven covariance matrices of dimension $p = 5$: three diagonal and four non-diagonal matrices. The first, A, is an identity matrix; the second, B, represents a covariance matrix with roughly equally spaced eigenvalues increasing in powers of 0.75 from 0.75^0 to 0.75^4 ; the third, C, is a somewhat ill-conditioned matrix, with eigenvalues equal to $0.75^0, 0.75^1, 0.75^2, 0.75^{10}$ and 0.75^{20} . They are then combined with rotations to produce four full true covariance matrices: B1 and C1 are matrix B and matrix C with Givens angles all set to $\frac{\pi}{4}$, B2 and C2 are matrixes B and C with Givens angles evenly spaced between $(-\frac{\pi}{4}, \frac{\pi}{4})$.

From each covariance matrix, we do the following simulation process:

- simulate a sample of size n ,
- compute $\hat{\Sigma}_{1,L_1}$, $\hat{\Sigma}_{2,L_1}$, $\hat{\Sigma}_{DK,L_1}$, $\hat{\Sigma}_{1,L_2}$, $\hat{\Sigma}_{2,L_2}$ and $\hat{\Sigma}_{DK,L_2}$ respectively the estimators for Σ from Model 1 of Section 3.2 (with $\delta = 2$), from Model 2 of Section 3.3 (with $\delta = 2$), from Model D&K of Section 3.1 (with $b = 10^6$), under L_1 and L_2 losses,
- compute $\hat{\Sigma}_{MLE}$,
- compute the associated loss $L_i(\hat{\Sigma}, \Sigma)$ for $i = 1, 2$ and for each estimator $\hat{\Sigma}$.

These estimations are carried out with the Metropolis-within-Gibbs algorithms described in Section 4.2 with 20,000 iterations, among which 5,000 are used for the burn-in period.

We then compare the different estimators with respect to the risk function $R_i(\hat{\Sigma}, \Sigma) = \mathbb{E}_{\Sigma} \left(L_i(\hat{\Sigma}, \Sigma) \right)$ $i = 1, 2$. These frequentist risks are approximated by repeating 100 times the previous simulation process.

From Propositions 1 and 2, δ must be strictly greater than 1. For the remainder of this paper, we will choose, quite arbitrarily, the smallest possible integer for δ , that is, $\delta = 2$. As an extension to this work, we could assign a prior on this parameter.

5.2. Performance comparisons

Here we proceed to comparison of risks - associated to a specific loss, L_1 or L_2 - between the prior models. Then we compare the ability to accurately estimate the eigenvalues across all competing estimators.

5.2.1. Under the Frobenius loss

Table 1 summarizes the simulation results for the frequentist risk R_2 for sample sizes $n = 5$ and $n = 100$.

The main remark is that, for $n = 5$, the shrinkage estimator from Model 1 always leads to (sometimes dramatic: risk divided by 2) improvement in accuracy over the alternative procedures. Secondly the estimators from Model 2 and Model D&K fail to accurately estimate and do even worse than the maximum likelihood estimator in all cases. The models in question lead only to shrinkage of the off-diagonal elements and such regularization methods tend to not have the expected beneficial effect of being more precise, under L_2 loss. In contrast Model 1 involves more severe shrinkage, which is given a very positive welcome. For a bigger sample size, the differences in performance between the estimators disappear except for the identity-matrix case (A).

5.2.2. Under Stein's loss

Table 2 gives the simulation results for the frequentist risk R_1 for sample sizes $n = 5$ and $n = 100$.

When $n=5$, the Bayes estimators provide substantial improvement in risk compared to the sample covariance matrix for the well-conditioned matrices (A, B, B1, B2), with risks from 3.78 to 8.7 times smaller. In these cases, Model 1 does somewhat better than Model 2 and Model D&K but as soon as Σ is ill-conditioned (C, C1, C2), it does worse than all other estimators. For the ill-conditioned diagonal matrix (C), the estimators from Model 2 and Model D&K perform well compared to the sample covariance matrix while they do poorly for its two rotated versions of matrices (C1, C2).

As expected, when the sample size becomes large (here $n=100$), the differences in performance between the estimators become blurred. Nevertheless in the well-conditioned cases the Bayes estimators still offer a non-negligible percentage reduction in risk compared to the sample covariance matrix.

	n = 5			
True matrices	$\hat{\Sigma}_{1,L_2}$	$\hat{\Sigma}_{2,L_2}$	$\hat{\Sigma}_{DK,L_2}$	$\hat{\Sigma}_{MLE}$
A	0.88 (0.13)	8.29 (0.71)	7.09 (0.59)	1.72 (0.12)
B	0.57 (0.06)	3.87 (0.45)	3.15 (0.39)	0.83 (0.08)
B1	0.55 (0.06)	3.75 (0.33)	3.22 (0.32)	0.79 (0.07)
B2	0.39 (0.04)	3.48 (0.30)	2.77 (0.23)	0.73 (0.06)
C	0.57 (0.05)	2.87 (0.46)	2.69 (0.78)	0.62 (0.09)
C1	0.45 (0.05)	2.39 (0.35)	2.39 (0.34)	0.60 (0.09)
C2	0.47 (0.06)	2.07 (0.27)	2.17 (0.28)	0.64 (0.08)
	n = 100			
A	0.03 (0.004)	0.11 (0.007)	0.10 (0.007)	0.09 (0.006)
B	0.04 (0.003)	0.04 (0.004)	0.04 (0.004)	0.04 (0.003)
B1	0.04 (0.003)	0.04 (0.004)	0.04 (0.004)	0.04 (0.004)
B2	0.04 (0.003)	0.05 (0.004)	0.04 (0.004)	0.04 (0.004)
C	0.04 (0.004)	0.04 (0.004)	0.04 (0.004)	0.04 (0.004)
C1	0.03 (0.003)	0.04 (0.004)	0.04 (0.004)	0.03 (0.003)
C2	0.03 (0.002)	0.03 (0.002)	0.03 (0.002)	0.03 (0.002)

Table 1: Comparison in risk between the Bayes estimators and the usual estimator of covariance matrix under L_2 . The values in parentheses refer to the simulation standard errors.

	n = 5			
True matrices	$\hat{\Sigma}_{1,L_1}$	$\hat{\Sigma}_{2,L_1}$	$\hat{\Sigma}_{DK,L_1}$	$\hat{\Sigma}_{MLE}$
A	0.66 (0.06)	1.42 (0.07)	1.18 (0.07)	5.75 (0.25)
B	0.87 (0.05)	1.50 (0.09)	1.26 (0.08)	5.61 (0.21)
B1	0.77 (0.05)	1.45 (0.07)	1.34 (0.07)	5.77 (0.23)
B2	0.70 (0.04)	1.47 (0.07)	1.34 (0.06)	5.52 (0.21)
C	42.42 (0.26)	1.37 (0.08)	1.17 (0.07)	5.69 (0.22)
C1	41.87 (0.26)	17.16 (0.97)	26.72 (2.39)	6.31 (0.32)
C2	42.77 (0.25)	26.36 (1.97)	42.04 (3.30)	6.22 (0.25)
	n = 100			
A	0.03 (0.003)	0.07 (0.003)	0.06 (0.003)	0.15 (0.006)
B	0.11 (0.004)	0.07 (0.004)	0.06 (0.004)	0.15 (0.005)
B1	0.12 (0.005)	0.13 (0.005)	0.13 (0.006)	0.16 (0.006)
B2	0.13 (0.005)	0.13 (0.005)	0.14 (0.006)	0.16 (0.006)
C	0.18 (0.007)	0.07 (0.004)	0.06 (0.003)	0.15 (0.005)
C1	0.18 (0.007)	0.16 (0.006)	0.16 (0.006)	0.16 (0.006)
C2	0.19 (0.006)	0.17 (0.006)	0.16 (0.006)	0.16 (0.006)

Table 2: Comparison between estimators with respect to the risk function R_1 . The values in parentheses refer to the simulation standard errors.

	$\hat{\Sigma}_{1,L_1}$	$\hat{\Sigma}_{2,L_1}$	$\hat{\Sigma}_{DK,L_1}$	$\hat{\Sigma}_{1,L_2}$	$\hat{\Sigma}_{2,L_2}$	$\hat{\Sigma}_{DK,L_2}$	$\hat{\Sigma}_{MLE}$
A	0.48 (0.02)	0.73 (0.02)	0.67 (0.02)	0.25 (0.02)	0.45 (0.02)	0.46 (0.02)	0.97 (0.01)
B	0.35 (0.03)	0.58 (0.02)	0.50 (0.03)	0.65 (0.06)	0.43 (0.03)	0.44 (0.03)	0.95 (0.01)
B1	0.29 (0.02)	0.52 (0.02)	0.45 (0.02)	0.71 (0.05)	0.50 (0.04)	0.49 (0.04)	0.95 (0.01)
B2	0.29 (0.02)	0.54 (0.02)	0.47 (0.02)	0.67 (0.05)	0.45 (0.04)	0.42 (0.03)	0.95 (0.01)
C	43.58 (2.62)	0.43 (0.03)	0.43 (0.04)	93.67 (4.47)	1.01 (0.09)	0.98 (0.10)	0.79 (0.02)
C1	43.15 (2.63)	11.47 (0.74)	13.96 (1.05)	94.28 (4.27)	36.87 (2.12)	40.38 (2.45)	0.86 (0.02)
C2	44.06 (2.50)	19.26 (1.10)	26.82 (1.89)	97.61 (4.64)	65.21 (3.48)	76.81 (4.59)	0.85 (0.02)

Table 3: Comparison in risk under $L_{\lambda_{min}}$, between the smallest eigenvalue of the competing estimators when $n = 5$. The values in parentheses refer to the standard errors.

5.2.3. About eigenvalues estimation

An important issue in covariance matrix estimation is the bias of the estimators of the extreme eigenvalues. We can then compare the bias in the eigenvalues of estimates based on the different models and approaches and we try to enlighten the previous results.

We set

$$L_{\lambda_i}(\lambda_i, \hat{\lambda}_i) = \frac{|\lambda_i - \hat{\lambda}_i|}{\lambda_i}, \text{ with } i \in \{min, max\} \quad (14)$$

as a relative measure of distance between the true (λ_i) and inferred eigenvalue ($\hat{\lambda}_i$).

Tables 4 and 3 give results for the minimal (λ_{min}) and maximal (λ_{max}) eigenvalues of the seven estimators for each type of covariance matrix. The eigenvalues of $\hat{\Sigma}_{MLE}$ differ greatly from the true values, especially when Σ is close to the identity matrix. $\hat{\Sigma}_{1,L_1}$, $\hat{\Sigma}_{2,L_1}$, $\hat{\Sigma}_{DK,L_1}$ and $\hat{\Sigma}_{1,L_2}$ appear to successfully estimate λ_{max} in all cases but fail in estimating λ_{min} in case of a strongly misspecified target. $\hat{\Sigma}_{2,L_2}$ and $\hat{\Sigma}_{DK,L_2}$ can improve λ_{min} but is always inappropriate for an accurate estimation of λ_{max} .

These results, combined with the boxplots of the smallest and largest eigenvalues, presented Annexe Appendix B, highlight that the Bayes methods reduce the distortion of the eigenvalue spectrum but that this effect can be too pronounced. Moreover the results stress that the performances under L_1 loss can be explained by the eigenvalues shrinkage: overshrinkage of the small eigenvalues will imply a poor performance under L_1 loss whereas it will not be penalized by L_2 loss which focuses on the errors on big values.

Estimator derived from Model 1 will be forced to be well-conditioned. Consequently it will yield to overshrinkage as soon as the true matrix has its eigenvalues far apart, but only of the smallest eigenvalues.

The overestimation phenomenon can be explained by the lower bound of the hyperparameter β as illustrated in Section 5.2.4.

5.2.4. Posterior distribution of β

As seen in Section 2, the hyperparameter β is known to control the amount of shrinkage. As the use of hierarchical models allows to estimate β from data,

	$\hat{\Sigma}_{1,L_1}$	$\hat{\Sigma}_{2,L_1}$	$\hat{\Sigma}_{DK,L_1}$	$\hat{\Sigma}_{1,L_2}$	$\hat{\Sigma}_{2,L_2}$	$\hat{\Sigma}_{DK,L_2}$	$\hat{\Sigma}_{MLE}$
A	0.27 (0.02)	0.59 (0.04)	0.64 (0.05)	0.84 (0.05)	2.26 (0.10)	1.99 (0.10)	1.70 (0.08)
B	0.34 (0.02)	0.37 (0.03)	0.40 (0.06)	0.33 (0.03)	1.21 (0.10)	1.02 (0.09)	0.68 (0.06)
B1	0.29 (0.02)	0.28 (0.02)	0.31 (0.03)	0.37 (0.04)	1.24 (0.08)	1.03 (0.08)	0.84 (0.07)
B2	0.29 (0.02)	0.27 (0.02)	0.29 (0.02)	0.30 (0.02)	1.14 (0.07)	0.94 (0.06)	0.71 (0.05)
C	0.41 (0.02)	0.33 (0.03)	0.34 (0.04)	0.38 (0.03)	1.10 (0.09)	0.96 (0.10)	0.57 (0.05)
C1	0.38 (0.02)	0.31 (0.02)	0.33 (0.03)	0.42 (0.04)	1.06 (0.08)	1.02 (0.09)	0.66 (0.06)
C2	0.37 (0.02)	0.31 (0.02)	0.34 (0.02)	0.40 (0.04)	0.93 (0.07)	0.90 (0.08)	0.67 (0.06)

Table 4: Comparison in risk under $L_{\lambda_{max}}$, between the largest eigenvalue of the competing estimators when $n = 5$. The values in parentheses refer to the simulation standard errors.




Figure 1: Histogram of the posterior samples of β from (11) when $n=100$ in the full ill-conditioned case (C2). The red line represents the lower bound of β .

data will determine the shrinkage intensity. Therefore, if the structured scale matrix is close to the true matrix, then β would take high values. Inversely one expects to get values for β arbitrarily small in the case where data support structure far from the prior scale matrix configuration. However, as mentioned by Daniel and Kass [3], a low intensity for shrinkage will be impossible to obtain because β must be always bounded by $p - 1$ minimum.

For illustration Figure 1 represents the histogram of posterior samples of β issued from Model 2 (a sample path of length 15,000 for one of the 100 datasets) when the true covariance matrix is C2 and $n = 100$. The posterior distribution of β is very concentrated on the lower bound $p+1$.

6. Conclusion and discussion

In this paper we proposed hierarchical Bayesian shrinkage methods for the estimation of covariance matrices in a small sample setting. Inverse-Wishart priors were considered with unknown hyperparameters. We focused on a "diagonal, common variance" and two "diagonal, unequal variance" models as covariance targets. Models are with 2 to $p+1$ free parameters, on which we assigned non-informative priors. We showed in details the conditions to ensure the property of the posterior distributions and proposed a Metropolis-Hasting-within-Gibbs algorithm to sample from them. Then we gave a detailed comparison between the different Bayesian estimators and the classical maximum likelihood covariance estimator, under three loss functions.

As statistically efficient and computationally fast alternative to the widely used standard covariance estimators, we recommend the shrinkage covariance estimators which shrink all components of the empirical covariance matrix, that is not only perfectly applicable to small samples but can also improve the classical estimators for large n . These improved estimators exhibit none of the defects of the standard covariance estimators, in particular they reduce variance, they are always positive definite and well-conditioned. This property might imply overestimation of the small eigenvalue. By producing a well-conditioned positive definite covariance estimate one automatically also obtains an equally well-conditioned estimate of the inverse covariance - a quantity of crucial importance, for instance, in classification or graphical models. For other goals like reduction dimension, principal component analysis needs to successfully estimate the largest eigenvalues rather than the smallest ones. However, for applications where we have to focus on estimation of all eigenvalues, care must be taken. Indeed under Stein's loss function, the proposed models can be unefficient. This happens when the data is in conflict with the specified prior structure, it is directly due to the lower bound of the degrees of freedom of an inverse-Wishart prior. Overshrinkage of the small eigenvalues is then caused.

The evidence on differences in estimation performance of different estimation approaches and models suggest that there is no "best approach" and that the relative accuracy of one approach or model in comparison to another depends strongly on the problem, that's one of the basic principles that underlie the Bayes paradigm.

A direct perspective of this work would be to investigate more complex models to overcome the issue of overshrinkage. Further work will be to investigate the effects of the estimation procedures in real data on the Value at Risk (VaR) computation. The VaR is a widely used tool for risk assessment in finance and is defined as a quantile of the predictive probability distribution for the amount of a future financial loss. The standard method for approximating the VaR is based on calculations using Monte-Carlo simulations of asset prices from a Gaussian distribution with unknown covariance matrix. The classical estimators have no full rank. Besides, in this application, focus is on the largest eigenvalues and then the proposed shrinkage estimator seems to be appropriate.

Acknowledgments

We wish to thank Christian Robert and Jean-Michel Marin for their contribution.

We would like to thank EDF (company of Electricity of France), especially Financial Direction, to support the thesis of Mathilde Bouriga.

References

- [1] Barnard, J., McCulloch, R., Meng, X., 2000. Modeling covariance matrices in terms of standard deviations and correlations, with applications to shrinkage. *Statistica Sinica* 10, 1281–1311.
- [2] Champion, C., 2001. Empirical Bayesian estimation of normal variances and covariances. *Journal of Multivariate Analysis* 26 (2), 60–79.
- [3] Daniels, M., Kass, R., 1999. Nonconjugate bayesian estimation of covariance matrices and its use in hierarchical models. *Journal of the American Statistical Association* 94 (448), 1254–1263.
- [4] Dempster, A., 1969. *Elements of Continuous Multivariate Analysis*. Addison-Wesley, Reading, Mass.
- [5] Dey, D., Srinivasan, C., 1985. Estimation of a covariance matrix under Stein’s loss. *The Annals of Statistics* 13 (4), 1581–1591.
- [6] Diaconis, P., Ylvisaker, D., 1979. Conjugate priors for exponential families. *The Annals of Statistics* 7 (2), 269–281.
- [7] Eaton, M., 1983. *Multivariate Statistics : A Vector Space Approach*. Wiley.
- [8] Efron, B., Morris, C., 1974. Multivariate empirical bayes and estimation of covariance matrices. *The Annals of Statistics* 4, 22–32.
- [9] Franck, C., Zakoïan, J., 2000. Covariance matrix estimation for estimators of mixing weak arma models. *Journal of statistical planning and inference* 83 (2), 369–394.
- [10] Furrer, R., Bengtsson, T., 2007. Estimation of high-dimensional prior and posteriori covariance matrices in kalman filter variants. *Journal of Multivariate Analysis* 98, 227–255.
- [11] Gelman, A., 2006. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* 1 (3), 515–533.
- [12] Haff, L., 1979. Estimation of the inverse covariance matrix : random mixtures of the inverse wishart matrix and the identity. *The Annals of Statistics* 7 (6), 1264–1276.
- [13] Haff, L., 1980. Empirical bayes estimation of the multivariate normal covariance matrix. *The Annals of Statistics* 8 (3), 586–597.
- [14] James, W., Stein, C., 1961. Estimation with quadratic loss. *Proceedings of the Fourth Berkeley Symposium on Mathematical and Statistical Probabilities Vol. 1*, University of California Press, Berkeley, 361–379.
- [15] Jeffreys, H., 1961. *Theory of probability*. 3rd ed. Oxford Classic Texts in the Physical Sciences, Oxford: Oxford University Press.

- [16] Kubokawa, T., 2004. A revisit to estimation of the precision matrix of the Wishart distribution. Unpublished discussion paper.
- [17] Kubokawa, T., Srivastava, M., 1999. Estimating the covariance matrix : a new approach. Unpublished discussion paper.
- [18] Ledoit, O., Wolf, M., 2002. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of empirical finance* 10, 603–621.
- [19] Leonard, T., Hsu, J., 1992. Bayesian inference for a covariance matrix. *The Annals of Statistics* 20 (4), 1669–1696.
- [20] Lin, S., Perlman, M., 1984. A monte carlo comparison of four estimators of covariance matrix. Tech. Rep. 44, University of Washington, Dept. of statistics.
- [21] Robert, C., 2001. *The Bayesian Choice: from Decision-Theoretic Motivations to Computational Implementation*. Springer.
- [22] Robert, C., Marin, J., 2007. *Bayesian core*. Springer-Verlag, New York.
- [23] Roberts, G., Rosenthal, J., 2006. Harris recurrence of metropolis-within-gibbs and trans-dimensional markov chains. *The Annals of Applied Statistics* 16 (4), 2123–2139.
- [24] Schäfer, J., Strimmer, K., 2005. A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Statistical applications in genetic and molecular biology* 4 (32).
- [25] Smith, M., Kohn, R., 2002. Parsimonious covariance matrix estimation for longitudinal data. *Journal of the American Statistical Association* 97 (460), 1140–1153.
- [26] Stein, C., 1956. Some problems in multivariate analysis. Tech. Rep. 6, Stanford University, Dept. of Statistics.
- [27] Stein, C., 1975. Estimation of a covariance matrix. Rietz lecture, 39th annual meeting IMS. Atlanta, Georgia.
- [28] Yang, R., Berger, J., 1994. Estimation of a covariance matrix using the reference prior. *The Annals of Statistics* 22 (3), 1195–1211.

Appendix A. Proof of Proposition 2

We want to evaluate whether

$$\int_{p+1}^{+\infty} \int_{]0;+\infty[^p} \int_{S^+} \pi(\Sigma, \Phi, \beta|S) d\Sigma d\phi_1 \dots d\phi_p d\beta \quad (\text{A.1})$$

is finite or not.

By Fubini's theorem:

$$\int_{p+1}^{+\infty} \int_{]0;+\infty[^p} \int_{S^+} \pi(\Sigma, \Phi, \beta|S) d\Sigma d\phi_1 \dots d\phi_p d\beta = \int_{p+1}^{+\infty} \pi(\beta|S) d\beta. \quad (\text{A.2})$$

Thus it is sufficient to prove the convergence of the posterior marginal distribution of β .

We begin to marginalize the density $\pi(\Sigma, \Phi, \beta|S)$ over Σ to obtain an expression for $\pi(\Phi, \beta|S)$. We have:

$$\pi(\Phi, \beta|S) \propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{|\Phi|^{\frac{\beta}{2}-1}}{|S + \Phi|^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta}.$$

Denote $Q\Lambda Q^T$, the spectral decomposition of the matrix S . Then

$$\begin{aligned} \pi(\Phi, \beta|S) &\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{|\Phi|^{\frac{\beta}{2}-1}}{|Q\Lambda Q^T + \Phi|^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta} \\ &\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{|\Phi|^{\frac{\beta}{2}-1}}{|Q\Lambda Q^T + Q\Phi Q^T|^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta} \\ &\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{|\Phi|^{\frac{\beta}{2}-1}}{|Q(\Lambda + \Phi)Q^T|^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta} \\ &\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{|\Phi|^{\frac{\beta}{2}-1}}{(|QQ^T| |\Lambda + \Phi|)^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta} \\ &\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{\prod_{i=1}^p \phi_i^{\frac{\beta}{2}-1}}{\prod_{i=1}^p (\lambda_i + \phi_i)^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta}. \end{aligned}$$

$$\text{Hence } \pi(\beta|S) \propto \int_{]0;+\infty[^p} \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{\prod_{i=1}^p \phi_i^{\frac{\beta}{2}-1}}{\prod_{i=1}^p (\lambda_i + \phi_i)^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta} \prod_{i=1}^p d\phi_i.$$

We just need to find a density $g(\beta)$ which dominates the positive function $\pi(\beta|S)$ on $]p+1;+\infty[$ and define the conditions so that the latter converges.

Let λ_{\min} be the smallest eigenvalue of S , it follows:

$$\begin{aligned} \pi(\beta|S) &\leq \int_{]0;+\infty[^p} \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{\prod_{i=1}^p \phi_i^{\frac{\beta}{2}-1}}{\prod_{i=1}^p (\lambda_{\min} + \phi_i)^{\frac{\beta+n}{2}}} \frac{1}{\beta^\delta} \prod_{i=1}^p d\phi_i \quad \forall \beta \in]p+1;+\infty[\\ &= \pi(\beta|S = \lambda_{\min} I). \end{aligned} \quad (\text{A.3})$$

Let us check if $\pi(\beta|S = \lambda_{min}I)$ is an integrable function over $]p+1; +\infty[$.

We have

$$\begin{aligned}
\int_{]0; +\infty[^p} \pi(\phi_1, \dots, \phi_p, \beta|S = \lambda_{min}I) \prod_{i=1}^p d\phi_i &\propto \int_{]0; +\infty[^p} \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \sqrt{\frac{\prod_{i=1}^p \phi_i^\beta}{\prod_{i=1}^p (\phi_i + \lambda_{min})^{\beta+n}}} \frac{1}{\prod_{i=1}^p \phi_i} \frac{1}{\beta^\delta} \prod_{i=1}^p d\phi_i \\
&\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{1}{\beta^\delta} \prod_{i=1}^p \int_0^{+\infty} \sqrt{\frac{\phi_i^\beta}{\left[\frac{\lambda_{min}}{n} \left(\frac{n}{\lambda_{min}} \phi_i + n\right)\right]^{\beta+n}}} \frac{1}{\phi_i} d\phi_i \\
&\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \frac{1}{\beta^\delta} \prod_{i=1}^p \int_0^{+\infty} \sqrt{\frac{\phi_i^\beta}{\left[\frac{\lambda_{min}}{n} \left(\beta \frac{n}{\lambda_{min} \beta} \phi_i + n\right)\right]^{\beta+n}}} \frac{1}{\phi_i} d\phi_i
\end{aligned}$$

In each single integral over ϕ_i , the density of a Fisher F-distribution with β and n degrees of freedom appears for the random variables $\frac{n\phi}{\lambda_{min}\beta}$. It implies:

$$\begin{aligned}
&\int_0^{+\infty} \sqrt{\frac{\left(\frac{n\phi_i}{\lambda_{min}}\right)^\beta n^n}{\left(\frac{n\phi_i}{\lambda_{min}} + n\right)^{n+\beta}}} \frac{\frac{n}{\beta\lambda_{min}}}{\phi_i} \text{Beta}\left(\frac{\beta}{2}, \frac{n}{2}\right) d\phi_i \\
&= \int_0^{+\infty} \sqrt{\frac{\phi_i^\beta \lambda_{min}^n}{(\phi_i + \lambda_{min})^{n+\beta}}} \frac{1}{\phi_i} \text{Beta}\left(\frac{\beta}{2}, \frac{n}{2}\right) d\phi_i = 1
\end{aligned}$$

and thus leads to

$$\begin{aligned}
\pi(\beta|S = \lambda_{min}I) &\propto \frac{\Gamma_p\left(\frac{\beta+n}{2}\right)}{\Gamma_p\left(\frac{\beta}{2}\right)} \prod_{i=1}^p \text{Beta}\left(\frac{\beta}{2}, \frac{n}{2}\right) \frac{1}{\beta^\delta} \\
&\propto \frac{\Gamma_p\left(\frac{\beta+n}{2}\right)}{\Gamma_p\left(\frac{\beta}{2}\right)} \frac{\Gamma\left(\frac{\beta}{2}\right)^p \Gamma\left(\frac{n}{2}\right)^p}{\Gamma\left(\frac{\beta+n}{2}\right)^p} \frac{1}{\beta^\delta}.
\end{aligned}$$

From the integration properties for positive functions, the inequality () implies the same inequality for their integrals. Hence

$$\int_{p+1}^{+\infty} \pi(\beta|S) \leq \int_{p+1}^{+\infty} \int_{]0; +\infty[^p} \pi(\phi_1, \dots, \phi_p, \beta|S = \lambda_{min}I) \prod_{i=1}^p d\phi_i d\beta.$$

This function is well-defined in $p+1$. Let us see the behaviour in $+\infty$.

By definition $\Gamma_p(\beta) = \pi^{\frac{p(p-1)}{4}} \prod_{j=1}^p \Gamma(\beta + \frac{1-j}{2})$. Furthermore Stirling's formula provides an approximation for the Gamma function: $\Gamma(\beta) \underset{\beta \rightarrow +\infty}{\sim} \exp^{-\beta} \beta^{\beta-1/2} (2\pi)^{1/2}$.

In consequence we get

$$\Gamma_p(\beta) \underset{\beta \rightarrow +\infty}{\sim} \pi^{\frac{p(p-1)}{4} + 2p} 2^{p/2} \exp\left(\sum_{j=1}^p \frac{j-1}{2}\right) \exp(-p\beta) \prod_{j=1}^p (\beta + \frac{1-j}{2})^{\beta - \frac{j}{2}}.$$

Hence

$$\begin{aligned}
\frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} &= \frac{\prod_{j=1}^p \Gamma(\frac{\beta+n+1-j}{2})}{\prod_{j=1}^p \Gamma(\frac{\beta+1-j}{2})} \\
&\stackrel{\sim}{\beta \rightarrow +\infty} \frac{\exp(-p\frac{\beta+n}{2}) \prod_{j=1}^p (\frac{\beta+n-1+j}{2})^{\frac{\beta+n-j}{2}}}{\exp(-p\frac{\beta}{2}) \prod_{j=1}^p (\frac{\beta-1+j}{2})^{\frac{\beta-j}{2}}} \\
&\stackrel{\sim}{\beta \rightarrow +\infty} \frac{\prod_{j=1}^p \beta^{\frac{\beta+n-j}{2}}}{2^{pn/2} \exp(\frac{pn}{2}) \prod_{j=1}^p \beta^{\frac{\beta-j}{2}}} \\
&\stackrel{\sim}{\beta \rightarrow +\infty} \frac{\beta^{\frac{pn}{2}}}{2^{pn/2} \exp(\frac{pn}{2})}.
\end{aligned}$$

Moreover

$$\begin{aligned}
\frac{\Gamma(\frac{\beta}{2})}{\Gamma(\frac{\beta+n}{2})} &\stackrel{\sim}{\beta \rightarrow +\infty} \frac{\exp(-\frac{\beta}{2})(\frac{\beta}{2})^{\frac{\beta-1}{2}}}{\exp(-\frac{(\beta+n)}{2}) \left(\frac{(\beta+n)}{2}\right)^{\frac{(\beta+n)-1}{2}}} \\
&\stackrel{\sim}{\beta \rightarrow +\infty} \frac{\exp(\frac{n}{2})}{\left(\frac{\beta}{2}\right)^{\frac{n}{2}}}.
\end{aligned}$$

As a result

$$\begin{aligned}
\pi(\beta|S = \lambda_{\min} I) &\propto \frac{\Gamma_p(\frac{\beta+n}{2})}{\Gamma_p(\frac{\beta}{2})} \left(\frac{\Gamma(\frac{\beta}{2})}{\Gamma(\frac{\beta+n}{2})} \right)^p \frac{1}{\beta^\delta} \\
&\stackrel{\sim}{\beta \rightarrow +\infty} \frac{\beta^{\frac{pn}{2}}}{2^{pn/2} \exp(\frac{pn}{2})} \frac{\exp(\frac{pn}{2})}{(\frac{p}{2}\beta)^{\frac{pn}{2}}} \frac{1}{\beta^\delta} \\
&\stackrel{\sim}{\beta \rightarrow +\infty} \frac{1}{\beta^\delta}.
\end{aligned}$$

Thus, by (A.4), $\pi(\beta|S)$ is integrable as soon as $\delta > 1$. Together with (A.2), the same conclusion holds for $\pi(\Sigma, \phi, \beta|S)$ so that (A.1) is finite.

Appendix B. Complement for simulation results

These boxplots show the shrinkage effects on the estimates for the extreme eigenvalues.




Figure B.2: Boxplots of the smallest (left) and largest (right) eigenvalues estimates when the true matrices are diagonal. The horizontal lines represent the true values of the eigenvalues. Top-to-bottom: A, the identity matrix - B, the well-conditioned diagonal matrix - C, the ill-conditioned diagonal matrix.

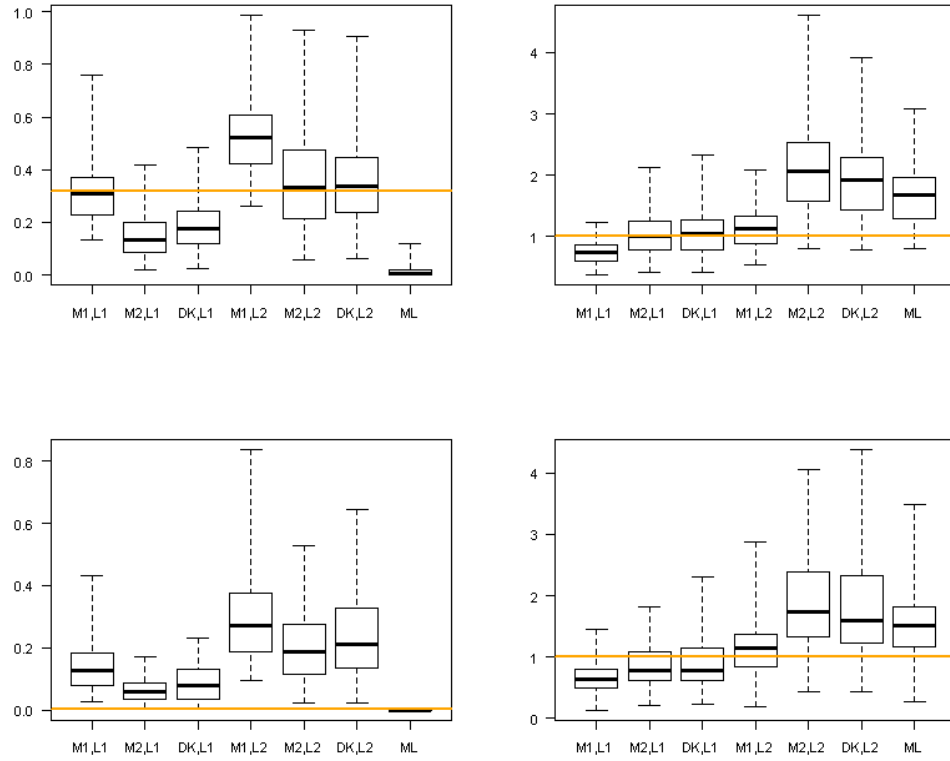


Figure B.3: Boxplots of the smallest (left) and largest (right) eigenvalues estimates when the true matrices are full. The horizontal lines represent the true values of the eigenvalues. Top-to-bottom: B2, the well-conditioned full matrix - C2, the ill-conditioned full matrix.